

Application No.: 09/827,362Docket No.: 30014165-2US (1509-164)Amendments to the Title:

Amend the title as follows:

METHOD OF ADDING A NEW USER ACCOUNT FROM A CLIENT COMPUTER  
ONTO A NETWORK INCLUDING HANDLING ON AGGREGATED GROUP OF  
MULTIPLE HEADLESS COMPUTER ENTITIES COMPUTERS AND NETWORK  
WHEREIN THE METHOD IS PERFORMED

Application No.: «SERIAL»Docket No.: «REFERENCE»**Amendments to the Specification:**

**Please replace the paragraph starting at line 22 on page 4 with the following amended paragraph:**

In the case of a system having a back-up computer entity for providing back-up data storage to a plurality of client computers, each ~~clients~~ client's back-up account is stored on a single computer entity, and this includes sharing common back-up data between accounts on that computer entity. In a best mode, an SQL database on the computer entity is used to keep track of the account data. This architecture means that the computer entities cannot be simply "clustered" together into a single logical entity. This means distributing the SQL database across all the computer entities in the group, and creating a distributive network file system for the data volumes across the computer entity group. This would be very difficult to implement, and it would mean that if one computer entity in the group failed, then the entire computer entity group would go off line.

**Please replace the paragraph starting at line 22 on page 17 with the following amended paragraph:**

Referring to Fig. 6 herein, there is illustrated schematically interaction of a plurality of headless computer entities 600-602, and a management ~~console~~ console(s) 617 in an aggregated group. The management console 617 comprises a web browser 604 which can view a web administration interface 605 on a master headless computer entity. The web interface on the master headless computer entity is used for some group configuration settings, including time zone setting and security settings. Other

Application No.: 09/827,362Docket No.: 30014165-2US (1509-164)

group administration function functions are provided by a Microsoft management console snap-in 616 provided on management console computer entity 617. Web interfaces 612, 613 are provided on each slave computer. The web administration interfaces on each computer entity are used to configure the computer entity level administration on those slave computer entities. On the [[mast]] master computer entity, the web administration interface 615 on that computer controls security and time zone settings for the entire group. All user application group level configuration settings are made via the MMC console 616 on the management console 617.

**Please replace the paragraph starting at line 7 on page 18 with the following amended paragraph:**

The master headless computer entity comprises an aggregation service application 607, which is a utility application for creating and managing an aggregation group of headless computer entities. The human operator configures a master user application 606 on the management console computer entity via the web administration interface 605 and web browser 604. Having configured the user application 606 on the master computer entity 600, via the management console 617, the aggregation service master application 607 keeps record of and applies those configuration settings across all slave headless computer entities 601, 602.

Application No.: 09/827,362

Docket No.: 30014165-2US (1509-164)

**Please replace the paragraph starting at line 17 on page 18 with the following amended paragraph:**

Each slave headless computer entity, 601, 602 is loaded with [[a]] the same aggregation service slave module 608, 609 and [[a]] the same slave user application 610, 611. Modifications to the configuration of the first application 606 of the master computer entity are automatically propagated by the aggregation service application 607 to all the slave applications 610, 611 on the slave computer entities.

**Please replace the paragraph starting at line 28 on page 18 with the following amended paragraph:**

Further, the master user application 606 on the master computer synchronises [[it's]] its application settings with each of the slave applications 610, 611 on the slave computers. The master user application 606 applies [[it's]] its synchronisation settings using the aggregation service provided by the aggregation service master and slave applications as a transmission platform for deployment of the user application settings between computer entities in the group.

**Please replace the paragraph starting at line 12 on page 19 with the following amended paragraph:**

Referring to Fig. 7 herein, there is illustrated logically an aggregation service provided by an aggregation service application 700, along with modes of usage of that service by one or more agents 701, data management application 702, and by a human administrator via web administration interface 703. In each case, the aggregation

Application No.: «SERIAL»Docket No.: «REFERENCE»

service master responds via a set of API calls, which interfaces with the operating system on the master headless computer entity. Operations are then propagated from the operating system on the master computer entity[[.]] to the operating systems on each of the slave headless computer entities, which, via the slave aggregation service applications 608, 609, make changes to the relevant slave applications on each of the slave computer entities.

**Please replace the paragraph starting at line 17 on page 20 with the following amended paragraph:**

Within the first group *Auto Back Up 1*, objects representing individual slave computer entities appear in sub groups including a first sub group consisting of protected computers, a second sub group consisting of users, and a third sub group consisting of appliance maintenance.

**Please replace the paragraph starting at line 25 on page 20 with the following amended paragraph:**

In the MMC-based management console, a menu option *create auto back up appliance group [[may]]* can be selected. This allows an administrator to create a computer entity group with the selected computer entity as the master. When creating the group, the administrator has the option to enable or disable an account balancing feature. The "account balancing" mode allows the administrator to provide the single

Application No.: 09/827,362

Docket No.: 30014165-2US (1509-164)

agent set up URL or agent download which automatically balances new accounts across the group.

**Please replace the paragraph starting at line 14 on page 22 with the following amended paragraph:**

In the simple case of Fig. 9, three individual groups each comprise three individual sets of computer entities, with no overlaps between groups. In the best mode herein, a single computer entity belongs only to one group, since the advantage of using the data processing and data storage capacity of a single computer entity is optimized by allocating the whole of that data processing capacity and data storage capacity to a single group. However, in other specific implementations and in general, a single computer entity [[may]] can serve in two separate groups, to improve efficiency of capacity usage of the computer entity, provided that there is no conflict in the requirements made by each group in terms of application configuration settings, or operating system configuration settings.

**Please replace the paragraph starting at line 26 on page 22 with the following amended paragraph:**

For example in a first group, a slave entity [[may]] can serve in the capacity of a network attached storage device. This entails setting configuration settings for a storage application resident on the slave computer entity to be controlled and regulated by a master computer entity mastering that group. However, the same slave computer entity [[may]] can serve in a second group for a different application, for example a graphics

Application No.: 09/827,362Docket No.: 30014165-2US (1509-164)

processing application, controlled by a second master computer entity, where the settings of the graphics processing application are set by the second master computer entity.

**Please replace the paragraph starting at line 10 on page 23 with the following amended paragraph:**

Referring to Fig. 10 herein, there is illustrated schematically actions taken by the aggregation service 607 when a new computer entity is successfully added to a group. The aggregation service 607 resident on the master computer entity 600 automatically synchronizes the security settings of each computer entity in the group in step 1001. This is achieved by sending a common set of security settings across the network, addressed to each slave entity within the group. When each slave entity receives those security settings, each slave computer entity self applies those security settings to itself. In step 1002, the aggregation service 607 synchronizes a set of time zone settings for the new appliance added to the group. Time zone settings will already exist on the master computer entity 600, (and on existing slave computer entities in the group). The time zone settings are sent to the new computer entity added to the group, which then applies those time zone settings on the slave aggregation service application in that slave computer entity, bringing the time zone settings of the newly added computer entity in line with those computer entities of the rest of the group. In step 1003, any global configuration settings for a common application in the group are sent to the client application on the newly added computer entity in the group. The newly added computer entity applies those global application configuration settings to

Application No.: «SERIAL»Docket No.: «REFERENCE»

the application running on that slave computer entity, bringing the settings of that client application[[,]] into line with the configuration settings of the server application and any other client applications within the rest of the group.

**Please replace the paragraph starting at line 11 on page 29 with the following amended paragraph:**

In step 1102, a third type of data management application configuration settings[[, are]] is applied such that any protected computer groups and their properties are synchronized across the group. The properties synchronized to the protected computer groups includes include schedule; retention; excludes; rights; limits and quotas; log critical files, and data file definitions applicable to protected computer groups. Again, this is effected by the master management application 606 applying those properties through the aggregation service 607 which sends data describing those properties to each of the computer entities within the group, which then self apply those properties to themselves.

**Please replace the paragraph starting at line 22 on page 30 with the following amended paragraph:**

The agent is created by an administrator using the MMC administration console, and an agent set-up utility available through that console. The agent is set-up on a computer entity selected from a list contained [[one]] on the master computer entity within a group. When the agent installs is installed on a slave computer entity, it will be automatically installed within a subgroup, and therefore will automatically pick up the

Application No.: 09/827,362Docket No.: 30014165-2US (1509-164)

policy settings of that subgroup. This requires [[that]] the master computer entity maintains to maintain a complete list of all subgroup settings for all subgroups within a group, that is, [[keeps]] to keep a list of which subgroups exist, and what the policy settings are for each of those subgroups, and the master computer entity synchronizes those subgroup policies across all slave computer entities within a group.

**Please replace the paragraph starting at line 12 on page 19 with the following amended paragraph:**

Therefore, for example once a computer entity is designated as a slave, it is no longer to perform subgroup management directly on that computer[[,]]; any subgroup management has to be done via the master computer entity in that subgroup. That is, one cannot access the slaved computer entity via the web interface on that slave directly[[,]]; any access must be through the master computer entity. The web administration interface effectively switches off some of the functionality at the level of the individual slave computer, and control can only be effected via the master computer entity for that grouped slave.

**Please replace the paragraph starting at line 19 on page 31 with the following amended paragraph:**

Referring to Fig. 15 herein, one way of using the systems disclosed herein, which may be beneficial to an external service provider, [[may be]] is as follows:

Application No.: «SERIAL»Docket No.: «REFERENCE»

**Please replace the paragraph starting at line 22 on page 31 with the following amended paragraph:**

Suppose the computer entity group is scaled up so that, for example, there are one million users of the computer entity group, the subgroup concept can be used to provide functionality for all [[a]] customers of the client computer entities, where the subgroup is tailored to the policies applied throughout all that companies computer entities.

**Please replace the paragraph starting at line 28 on page 31 with the following amended paragraph:**

For example, suppose a client company requires has five thousand users, an operator of a computer group system may create creates a subgroup for that company (step 1500), supplying five thousand users[[.]] with that companies company's particular policy settings applied across all slave computer entities within the subgroup (step [[1500]] 1501). The operator then gives the company an agent download (step 1502). When the company installs the agent onto all of [[their]] the computers of the company in step 1503, [[they]] the computers automatically pick all that companies company's policy settings from the agent[[.]]. As a result, the companies company's client computers are automatically capacity balanced across the slave computer entities in the subgroup operated by the operator, and the administrator in the external service provider has very little administration to do.

Application No.: «SERIAL»Docket No.: «REFERENCE»

**Please replace the paragraph starting at line 13 on page 32 with the following amended paragraph:**

Referring to Fig. 16 herein, there is illustrated schematically process steps carried out by the executable agent installation program and the master computer entity for initiating installation of an agent onto a slave computer entity. In process 1600, the executable agent installation program, having been received by a computer entity within a network, locates the master computer entity on the network. In step [[1501]] 1601, the executable agent installation program seeks instructions from the master computer entity as to [[which]] the slave computer entity to install the agent on which the agent is to be installed. In step 1602, the master entity queries all slave computer entities within the group, and determines which slave computer is best for installation of a new user account. Determination is based upon the parameters of firstly data storage capacity of the slave computer entity (step 1603), and secondly a sub-net mask of each of the slave compute entities. A local area network can be logically divided into a plurality of logical sub-networks. Each sub-network has its own range of addresses. Different sub-networks within the network are connected together by a router. The master computer entity attempts to match the sub-net on which a client computer for which an account is to be opened, with a slave back-up computer which is to provide that account (step 1604), so that the slave back-up computer and the client computer are both on the same sub-network, thereby avoiding having to pass data through a router between different sub-networks to provide the user account back-up service. Following step 1604, the master computer sends the identification of the slave computer on which the new user account is to be installed to the executable agent installation program, and the

Application No.: 09/827,362Docket No.: 30014165-2US (1509-164)

executable agent installation program proceeds to install the new user account on that specified slave computer.

**Please replace the paragraph starting at line 1 on page 34 with the following amended paragraph:**

If the computer entity group is in a "generic" security mode, then the agent set up web wizard ~~should check~~ checks that the new account being created by the user is unique across the entire computer entity group. This means that all the computer entities in the group must be on-line if the user tries to create a new account, and if any computer entity in the group is off-line then the user ~~should get~~ gets an error message in the agent set up wizard telling them this. If the new account is unique across the group, then the user downloads an AgentSetup.exe file, and runs this on their client computer.

**Please replace the paragraph starting at line 25 on page 34 with the following amended paragraph:**

All of the computer entities in the aggregated group need to be on-line in order to create a new account, because the account uniqueness check run by the AgentSetup.exe must be run across all the accounts in the entire appliance group. If one or more appliances in the group are off-line, then an error message ~~should be is~~ displayed by AgentSetup.exe telling the user that they cannot create their new back-up account.

Application No.: 09/827,362Docket No.: 30014165-2US (1509-164)

**Please replace the paragraph starting at line 8 on page 35 with the following amended paragraph:**

It is possible for computer entities within a group to have different sub-net settings. This is used for sites where there are multiple sub-nets and the computer entities within the group are configured on different sub-nets so that the backup traffic is kept within the sub-nets. Given this, the account balancing algorithm ~~needs to attempt~~ attempts to create the new account on a computer entity in the group which matches the client's sub-net mask. The algorithm to select which computer entity in the group should hold the new account restricts itself to just those computer entities which are valid targets and which have the same sub-net mask as the client. If there are no computer entities within the group which are valid targets, and which match the client's sub-net mask, then the algorithm selects any valid appliance target in the group to hold the new account, regardless of sub-net masks.

**Please replace the paragraph starting at line 21 on page 35 with the following amended paragraph:**

From the set of valid computer entity targets which match the client's sub-net mask, or all valid computer entity targets if there is no match, the algorithm selects a computer entity with the maximum available free data storage space compared with the other valid computer entities. If there are multiple computer entities with the same maximum available free space, then the algorithm randomly selects one of these. The agent set up procedure is then automatically and transparently redirected to the

Application No.: 09/827,362Docket No.: 30014165-2US (1509-164)

selected computer entity. After redirection, the agent set up runs to completion as normal, targeting the selected computer entity.

**Please replace the paragraph starting at line 6 on page 36 with the following amended paragraph:**

The account selection list shown during reinstallation of an existing account should be is a super set of the existing accounts on all the computer entities in the group. If the selected account is on a different computer entity in the group, then AgentSet.exe will automatically and transparently be redirected to continue the agent set up wizard on that computer entity. This therefore requires the master computer entity in the group to be online in order to provide a list of group members, and thus query all the computer entities in the group for the list of current accounts. If any slave computer entities are off line when the master runs the query, then any accounts held in the off line slave computer entities will not be displayed in the reinstallation account list. However, if the master computer entity is off line and the agent set up is run from one of the slave computer entities, then only the accounts held on that slave computer entity will be displayed.

**Please replace the paragraph starting at line 6 on page 37 with the following amended paragraph:**

If a data management application uses the aggregation features of the aggregation service application 700, [[then]] and if the management application is user or account centric, in the best mode an account balancing scheme is used across the computer

Application No.: 09/827,362Docket No.: 30014165-2US (1509-164)

entity group. The aim of this account balancing is to treat the computer entity group as a single logical entity when creating new data management application accounts, which means that an administrator does not have to allocate individual accounts to specific computer entities.

**Please replace the paragraph starting at line 14 on page 37 with the following amended paragraph:**

For example, when creating a new account, the data management application may obtain obtains a current computer entity group structure using a *read appliance group structure API*, ~~and then use~~. The data management application then uses this information to query each data management application on every computer entity in the group. The account can then be installed on the computer entity in the group which best meets the data management application criteria for a new account, for example the computer entity with the most free data storage capacity available.

**Please replace the paragraph starting at line 22 on page 37 with the following amended paragraph:**

If the data management application does implement account balancing, then the administration should have the option, when creating a computer entity group, to enable or disable this mode. It is possible to disable the account balancing mode for the cases where the administrator ~~wanted~~ wants to be able to create a computer entity group across multiple different geographic sites for the purposes of setting data management

Application No.: «SERIAL»Docket No.: «REFERENCE»

policies. However, in this case, the administrator would want to keep the accounts for one site on the computer entities on that site, due to the network traffic.

**Please replace the paragraph starting at line 1 on page 38 with the following amended paragraph:**

Referring to Fig. 17, there is illustrated schematically a network of a plurality of computer entities, comprising: a plurality of client computer entities  $C_1 - C_N, C_{N+1} - C_{N+M}$ , each client computer entity typically comprising a data processor, local data storage, memory, communications port, and user console having a visual display unit, keyboard and pointing device, e.g. mouse; a plurality of headless computer entities, the headless computer entities designated as master computer entities  $M_1, M_2$ , and slave computer entities,  $S_1 - S_6$ . The master and slave computer entities provide a service to the client computers, for example a back-up facility. The master and slave headless computer entities may comprise for example network attached storage devices (NAS). The plurality of computer entities are deployed on the network across a plurality of sub-networks, in [[he]] the example shown of first sub-network [[1600]] and a second sub-network [[1601]]. The two sub-networks, comprising the complete network, are connected via a router [[1602]] 1702. The headless computer entities are aggregated into groups, comprising a master computer entity and at least one slave computer entity. In the best mode implementation, computer entity groups are all contained within a same sub-network, although in the general case, an aggregation group of headless computer entities [[may]] can extend over two or more different sub-networks within [[a]] the same network.

Application No.: 09/827,362Docket No.: 30014165-2US (1509-164)

Please replace the paragraph starting at line 21 on page 38 with the following amended paragraph:

Referring to Fig. 18 herein, there is illustrated schematically process steps carried out by an account balancing algorithm for the process [[1700]] 1800 of setting up a new user account on a computer entity within an aggregated group. In step 1801, the algorithm checks that all computer entities within the group are on-line. If not, then in step 1802, the algorithm cannot create a new back-up account and in step 1803 displays an error message to a client computer that a new back-up account cannot be created. If however all computer entities within the group are on-line, then in step 1804 the algorithm runs an account uniqueness check amongst all the computer entities within the group. In step 1805, the algorithm identifies which computers in the group are valid targets to hold a new user account. Details of the operations of step 1805 are illustrated in Fig. 19, described infra. If no valid targets are found in step 1806, then the algorithm cannot create a new back-up account and displays an error message in step 1803 as described ~~in step 1803~~ previously. However, provided valid targets are found, then in step 1807 the algorithm compares a sub-net address of the client computer for whom the back-up account is to be created, with the sub-net addresses of all the valid targets found in the group. If valid targets target computers are found with the same sub-net address as the client computer in step 1808, then in step 1809 the valid target computers having [[a]] the same sub-net address as the client computer are selected to form a set of valid target computers in step 1811. However, if no valid target computers have [[a]] the same sub-net address as the client computer to which a user account is to

Application No.: «SERIAL»Docket No.: «REFERENCE»

be supplied, then in step 1810, the algorithm selects a set of all valid target computers within the same group, regardless of the sub-net mask, to form a set of valid target computers 1811.

**Please replace the paragraph starting at line 14 on page 39 with the following amended paragraph:**

In step 1812, the algorithm selects a valid target computer having [[a]] the maximum available free data storage space. If a computer entity having [[a]] the maximum available free data storage space cannot be selected in step 1813, for example because two computers have [[a]] the same amount of available free data storage space and no valid target computer has a maximum, then in step 1814 the algorithm randomly selects one of the valid target computers in the set 1811. In step 1815, the AgentSetup.exe is redirected to the selected target computer. In step 1816, the AgentSetup.exe program is run to completion, targeting the selected target computer, thereby creating one or more new accounts on that target computer for use by the client computer.

**Please replace the paragraph starting at line 25 on page 39 with the following amended paragraph:**

Referring to Fig. [[18]] 19 herein, there is illustrated schematically one implementation of process steps carried out in step 1805 to identify which computers in a group are valid targets to hold a new user account. In step 1900, a "next" target computer within a group is identified by the algorithm. In step 1901, the algorithm

Application No.: 09/827,362Docket No.: 30014165-2US (1509-164)

checks whether the "next" target computer has any available data storage space left. If the data storage space of the "next" target computer is full, then in step 1904 the algorithm identifies the "next" target computer as an invalid target computer. However, if the data storage space is not full, then in step 1902 the algorithm checks if the "next" target computer entity has reached a "*new user capacity limit*", [[being]] i.e., a limit at which new users cannot be taken onto [[that]] the "next" target computer entity. If that limit is reached, then [[that]] the "next" target computer is identified as an invalid target computer in step [[1804]] 1904. However, if the *new user capacity limit* has not been reached, then in step 1903, the "next" target computer is added to the list of available valid target computers. In step 1905, ~~it is checked a check is made as to~~ whether all possible target computers have been checked, and if not steps 1900 - 1905 are repeated until all target computers have been checked for validity.

**Please replace the paragraph starting at line 24 on page 40 with the following amended paragraph:**

Referring to [[fig.]] Fig. 20 herein, there is an illustrated schematically processes carried out by the management console MMC application 616 for automatic balancing of user accounts across the plurality of computers in a group. The MMC application 616 contains a user account migration component, which operates to move complete user accounts, [[eg.]] e.g., a client's backup account, from one computer entity within a group to another, without any impact on the user who owns that account.

Application No.: 09/827,362Docket No.: 30014165-2US (1509-164)

**Please replace the paragraph starting at line 1 on page 41 with the following amended paragraph:**

Since each user account is stored on a single computer entity within the group, if the data storage space on that computer entity becomes fully utilised, then there is no further capacity on that computer entity for addition of further data in the user account. Therefore, the user accounts must be moved from the "full" computer entity onto an "empty" computer entity. The empty computer entity can be any computer entity within the group having enough spare storage capacity to accommodate a user account moved from the full computer entity. An empty computer entity [[may]] can be a new computer entity added into the group, having un-utilised data storage capacity.

**Please replace the paragraph starting at line 24 on page 41 with the following amended paragraph:**

In step 2000, the MMC application monitors the utilised capacity on each of the plurality of computers in a group. The MMC application monitors the data storage capacity utilisation, and compares this with [[the]] hard and soft quota limits in each computer in the group. In step 2001, having found a computer entity is found having [[where]] utilisation of data storage space [[is]] above the soft quota limit, this. This indicates that that computer entity is becoming "full", [[ie.]] i.e., the data storage capacity is almost fully utilised. Therefore the MMC application continues in step 2002 to locate a "empty" computer entity within the group having enough free capacity to hold some accounts from the located full computer. The MMC console checks a "new user capacity limit" on each computer entity in the group, being a capacity limit for addition of

Application No.: 09/827,362Docket No.: 30014165-2US (1509-164)

a number of new users. If a suitable computer entity having a number of users below the new user capacity limit is not found in step 2003, then the MMC application 616 generates an alert message to the administrator to add a new computer to the group in step 2004. However, if the MMC application finds a suitable computer having a number of users below the new user capacity limit, then in step 2005, the MMC application selects user accounts from the full computer for relocation to the selected empty computer. Where more than one empty computer is found, the MMC application may select the empty computer randomly, or on the basis of lowest utilised capacity. In step 2006, once the master computer entity has determined which user accounts are to be transferred from the full computer to the selected empty computer or computers, to receive those user accounts, then the master computer configures and then initiates a user account migration job on the full computer. From this point, user account migration runs as though the administrator has manually configured a user account transfer using the MMC console. However, the process is initiated automatically without human administrator intervention. Therefore, even if all computers in the computer group are nearing full capacity a human administrator would only have to install a new empty slave computer into the group, and the automatic capacity leveling provided by the process of Fig. 20 would automatically start transferring accounts from full computers onto the newly added computer entity, so that capacity was freed up on the full computers in the group.